Wright, Joss

Article

Necessary and inherent limits to internet surveillance

Internet Policy Review

Provided in Cooperation with:
Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin

Suggested Citation: Wright, Joss (2013) : Necessary and inherent limits to internet surveillance, Internet Policy Review, ISSN 2197-6775, Alexander von Humboldt Institute for Internet and Society, Berlin, Vol. 2, Iss. 3, pp. 1-11, https://doi.org/10.14763/2013.3.184

This Version is available at:
https://hdl.handle.net/10419/213969

INTERNET POLICY REVIEW
Journal on internet regulation

Volume 2 | Issue 3

Necessary and inherent limits to internet surveillance

Joss Wright
Oxford Internet Institute, United Kingdom, joss.wright@oii.ox.ac.uk

Published on 05 Aug 2013 | DOI: 10.14763/2013.3.184

Abstract: Information technologies now play a huge role in both personal and institutional life, playing the role of a global communications medium. As our means of interaction increasingly centre on the internet, there is a desire from nation states to exercise control and obtain access to the communications of citizens. The stated reasons for this access and control are to prevent or investigate crimes, and to protect national security. This article argues that mass untargeted surveillance of internet-based communications is an excessive tool with respect to its potential for abuse against both society and individuals, and that its ability to prevent crime or terrorism are limited. By looking at existing technologies and example cases where surveillance has been applied, this article demonstrates that there are both inherent mathematical and technical limits to the potential for surveillance to achieve broad-scale prevention of crime and terrorism. In addition, the potential of surveillance to result in real harm to society necessarily places severe limits on how this technique should be applied in a free and democratic society.

Keywords: Surveillance, Intelligence services, Base rate fallacy, Interception, Fundamental principles, Predictive algorithms, Privacy, Security, Content data, Communications data, Bundestrojaner, National Security Agency (NSA), PRISM, Tempora, Encryption, Liability, Transparency, Cyber security, Content, Censorship, Filtering

Article information

Received: 30 May 2013 **Reviewed:** 24 Jul 2013 **Published:** 05 Aug 2013

Competing interests: The author has declared that no competing interests exist that have influenced the text.

URL: http://policyreview.info/articles/analysis/necessary-and-inherent-limits-internet-surveillance

Citation: Wright, J. (2013). Necessary and inherent limits to internet surveillance. *Internet Policy Review*, *2*(3). DOI: 10.14763/2013.3.184

1 OVERVIEW

Information technologies, and in particular the internet, have brought about fundamental changes in how our society functions. Perhaps the most fundamental of these changes is in the ways in which we communicate; whilst the ability of computers to store and process data has expanded rapidly it is, arguably, the rise of instant, global data transfer that has had the most

far-reaching effects.

E-mail, instant messaging, and peer-to-peer file transfers, combined with the digitisation of content, have changed how we experience the world, the means by which we access information, and the shape of our social networks.

The general-purpose nature of computing and telecommunications has, necessarily, resulted in applications of these technologies that are undesirable, illegal and socially unacceptable. There are crimes that are unique to the internet, such as hacking or distributed denial of service attacks against websites, but in many cases the internet simply provides a new medium for more traditional crimes: blackmail, fraud, or dealing in stolen property such as credit cards.

In light of the move to internet-based communications, that have largely replaced many traditional services, police and intelligence services are understandably concerned that criminal activities that take place on, or make use of, the internet should be subject to investigation and punishment. This article argues that there are significant dangers in surveilling online communications unless the mechanisms and policies of surveillance are subject to strict and legally enforceable standards of transparency, oversight, and control, both nationally and internationally.

Recent revelations regarding the pervasive and wide-ranging surveillance of internet communications by intelligence services, in particular those of the United Kingdom and the United States, highlights the power of the internet as a tool for the monitoring of individuals globally. At the same time, these revelations have demonstrated the ease with which surveillance policies can extend to a large proportion of the world population, and the difficulties of providing meaningful levels of oversight and transparency for these policies.

The technical capabilities of the internet not only allow for surveillance, they encourage us, through convenience, to place more and more of our lives into the spotlight. We now read news, search for information, talk to friends, organise social and business life, bank, and meet potential partners via the internet. There is no precedent that can even approximate a model for the pervasiveness of the internet in our lives – not the phone network, not post or telegraph, not CCTV surveillance. Equating the internet with historical technologies when making policy is not simply wrong, it is dangerously misleading.

In making use of the internet to underpin so many aspects of our lives, unprecedented levels of data can now be collected, stored, and analysed, and are increasingly combined and controlled in a largely centralised manner. While sharing this data is, in some cases, contentious, we are increasingly unable to interact with internet services without explicitly or implicitly revealing personal information. The algorithms used to infer future behaviour from this data have improved; the computers that run these algorithms have become faster and more able to handle larger volumes of data; and the results of these inferences are being used to drive decisions, from customising search results to offering, or refusing to offer, tailored products and services.

Despite this, there are both ethical and technological limits to data-driven surveillance. Simply because data is generated and can be stored does not suggest that states should abandon fundamental principles and surveil entire populations rather than targeted individuals (United Nations Human Rights Council, 2013). The improvements in predictive algorithms over recent years are still subject to fundamental limitations in their accuracy and applicability.

From the state's perspective, the desire for surveillance is easy to understand. Such a wealth of

data seems to promise an oracle allowing security services not only to investigate, but also to detect, predict and prevent crimes – and ubiquitous surveillance can, certainly, achieve some of these goals. From the perspective of citizens the balance between an abstract notion of privacy and a more direct feeling of security, whether actual or perceived, can lead to popular support for invasive surveillance measures (YouGov Poll, 2012).

The wealth of data that surveillance reveals, however, tips the balance decisively from its power to help towards its power to harm. Vast amounts of information can be handled by faster and faster computers, but the power and accuracy of predictive algorithms are not so scalable – when applied blindly to entire populations the ability to identify suspicious patterns is lost in the flood and becomes either worthless or actively harmful.

Pervasive and detailed information on individuals is a powerful tool. When investigating a crime the details of suspects' activities, communications, and habits can be highly valuable. This tool, however, can be used just as effectively against all those individuals who are not under suspicion – blackmail, fraud, stalking, and simple invasion of privacy are all enabled by such collections of data just as effectively as the investigation of crime. Whether such abuse comes from institutional malpractice, malicious insiders, or external attackers who gain access to stored data, it is the gathering of such data that enables abuse.

2 LAW AND TECHNOLOGY

There is a significant disparity between the rate at which technology develops and the rate at which new laws can be brought into force. The application of outdated laws to technologies that were inconceivable when the laws were drafted is now a common occurrence. The United Kingdom's Computer Misuse Act was passed in 1990, in response to a number of high-profile hacking incidents that took place in 1988, the year before the world wide web itself was first proposed. Even with significant amendments made fourteen years later this legislation, with no mention of continuously connected smartphones reporting our location and data and services moving to cloud computing, already appear worryingly outdated.

The inability of the law to keep pace with technology can lead to significant difficulties in appropriate interpretation of laws. The ongoing debates concerning the separation between metadata and content of communications are one example, but there have been significant discussion regarding when emails are considered to have been delivered, rather than being in transit, as well as the specifics of whether viewing an image on a website is considered to be 'viewing' or 'creating' an image[1].

It would arguably be far more dangerous for the law to attempt to track new developments too rapidly. A common principle in software engineering, particularly in the world of open source software, is summarised by the phrase 'release early, release often'; imperfect software is often made available as soon as it is even partially functional, allowing bugs to be discovered and fixed by the community. Laws, thankfully, are not typically created and tested in the same way. There is clearly a fine balance to be struck between applying increasingly out of date laws to new technologies, and in overly reactionary lawmaking. This is, however, an ongoing and dynamic process that neither has, nor is amenable to, a fixed solution.

Surveillance in particular, however, presents a particular concern within this landscape. Not only do the details of the technology have serious and subtle effects on the efficacy of the laws,

the laws themselves are often driven by emotive and high-profile events such as terrorism or child abuse.

Where the discrepancy between the pace of technological change and the pace of legal change can have severe effects, such as the development of surveillance tools and methods far beyond the scope considered possible when laws are drafted, lawmakers must consider carefully the risks that arise from the future development and application of technologies. Where limitations to the power of a surveillance approach are inherent in the current state of technology, legal restrictions may be absent. As technology develops, the technological limitations erode, which can lead to the extension of the surveillance far beyond its original intent.

The erosion of implicit limitations is exemplified in the ongoing debates, both as part of the NSA PRISM revelations and in rhetoric related to surveillance programmes such as the UK's proposed Communications Data Bill, regarding the difference between content data and communications data. A traditional view has held that communications data, or metadata, about communications without the content of those communications, is relatively harmless and so subject to much lower protections than content data. However the expanded range of services for which metadata are available, the uncertain delineation between content and communications data, and the improved capacity to store and analyse such data, have greatly increased the power of metadata analysis.

Crucially, and challengingly, it is necessary to differentiate between short-term limitations that exist in current technologies and that will disappear as technology develops, and those limitations that are fixed and inherent.

3 RISKS AND BENEFITS

The use of surveillance to investigate and prevent crime almost unavoidably carries with it the risk of infringing on individual rights to privacy and freedom of expression, such as those set out in the European Convention on Human Rights. It is crucial to consider both the benefits of the proposed approach in terms of achieving its stated goals, against the risks of failure, abuse and misapplication.

The benefits of surveillance technologies are easily expressed – an improved ability to detect, investigate, and prevent crime. By gathering more data, it is argued, crime can be reduced. While this notion is intuitively appealing, such a sweeping generalisation hides complexities. What types of crime are prevented by increased surveillance? How effective is surveillance in solving crimes, and how does this effectiveness change as levels of surveillance increase? These factors have been particularly studied (PDF) in the UK context of CCTV coverage; evidence has shown that the benefits are variable, and the effectiveness low when compared with the costs (Armitage, 2002). It is important both to understand the limitations of an increase in surveillance and to balance it against the risks to society of overly-invasive surveillance practices.

When investigating crime, the ability of police and intelligence services to gain information regarding a suspect's communications and activities is clearly a great benefit. The parties with whom an individual has been communicating, the frequency and duration of those communications, the content, the websites accessed, all provide significant insight. To a large extent, the ability to gain such information concerning a suspect is already provided for in laws

such as the UK's Regulation of Investigatory Powers Act. This information is not, generally, available for a suspect's past beyond the provisions of the EU Data Retention Directive, which is restricted largely to user account information and the email account provided by the internet service provider (ISP). There are, however, significant and growing powers to obtain data of individuals that have come to the attention of the relevant bodies, such as those set out in the French LOPPSI2 and German use of trojan horse programmes to gain access to the computers of suspects (Chaos Computer Club, 2011).

A separate issue is that of using blanket surveillance, as exemplified by the US PRISM and XKeyscore or UK Tempora programmes to detect and prevent crime. Intuitively, a similar argument applies: by data mining large volumes of population data, indicators of criminal activity can be isolated and the suspicious parties can be placed under closer surveillance. This argument, however, highlights an inherent flaw in the ability to detect and identify such patterns in large populations, and thus suggests that such an approach is not as beneficial as it might first appear.

When applied over large scale populations, a well-known statistical effect known as the base rate fallacy highlights that even extremely accurate detection methods quickly become unable to identify persons of interest meaningfully. This effect is due to high rates of false positives, in which an innocent individual is identified as suspicious; and false negative, in which a suspicious individual is not detected. A full discussion of this issue is outside of the scope of this article; see this link for a useful summary of the problem (Bar-Hillel, 1980).

We therefore see a significant distinction between the power of increased surveillance to be of use in investigating crime, and in detecting crime. For the former, increased access to information is likely to be of use; in the latter case there is a strong argument to be made that too much information over too great a population is counterproductive. Even for the investigation of crime, however, the risks of the misapplication of technology, and the implications of creating an infrastructure that allows for the gathering of such information, must be seriously considered.

4 ELEMENTS OF SURVEILLANCE

The internet is a complex and partially decentralised network of networks that is largely dissimilar to older technologies such as the telephone network. The services that run on this network – websites, email, voice calls, instant messaging, peer-to-peer networking – present a range of challenges for interception and analysis. This is further complicated by the increasing use of strong encryption as a default for online services.

There are, therefore, two major sources of data concerning the activities of internet users: data stored by service providers, such as Google and Facebook as well as ISPs; and data stored on devices, such as smartphones and computers, of users themselves. Both of these sources of information are targeted by online surveillance, and present their own difficulties and risks.

The European Data Retention Directive (2006/24/EC) aims towards the first of these, requiring communication service providers to retain information regarding the source, destination, and duration of certain classes of telecommunications. At present, however, this does not extend to key modern web-based services for communications, such as messages sent via Facebook, that pass through a third party.

A similar approach has been implied by the discussions surrounding the United Kingdom's Communications Data Bill. The proposed legislation sought to provide a means for law enforcement and intelligence services to extend monitoring of communications data to services such as GMail and Facebook, although the specific details of how this would have been achieved at a technical level were not revealed at the time. It was hypothesised that, due to the use of encryption, this would require direct cooperation of the service operators; this view has been supported by the revelations surrounding PRISM and Tempora.

What is clear, however, is that surveillance on the modern internet requires interaction with private corporations, who use strong encryption to protect their users' communications from hackers. With strong encryption on transfers, supported by strong authentication of the services to which users connect, the only realistic option[2] is requiring companies, including Facebook and Google but also smaller providers of similar services, to install infrastructure in order to perform interception. We must ask ourselves who is to purchase, install, control, and maintain this infrastructure? With whom does liability for misuse or misapplication of these technologies lie, and how can such a system provide transparency and safeguards against abuse? The recent details concerning these programmes has shown that there are few satisfactory answers in existing approaches.

A second approach targets user devices directly. This is an extremely powerful approach, as user devices are both a source and recipient of the users' communications and provide a single point of access for all services that a user accesses, as well as potentially providing further access to online services. Direct access to user devices is, of course, a difficult technical challenge and lends itself less readily to mass surveillance; software must be installed on user devices, presumably without their knowledge, which is infeasible for an entire population. This approach has been taken in Germany, as we shall see in §5.1, as well as being provided for in recent French security legislation.

What is perhaps a greater concern than the specific approach taken to surveillance, is that of the infrastructure that supports that approach. As has been discussed above, surveillance based in the service provider requires the installation of physical devices that allow access to communications, which must either be directly controlled by the state, or which are managed by the provider in response to requests from the state. In either case, a technical capability has been created to log and make available large amounts of user data, which in turn creates the potential for such data to be accessed without the correct authorisation, either by company employees, by members of the intelligence services without correct authorisation, or by third parties that gain access.

In the second case, which targets user devices, appropriate invasive software must be developed, along with the means for deploying it on user devices. As such software must be installed without user knowledge, there must either be provision for physical access to the device or for remote compromise of the machine via the internet. Again, with such software and procedures in place the potential for abuse is great.

Most notably, where police and intelligence services are directly engaged in hacking into remote machines, as has been proposed in France, Germany, and The Netherlands, the means to do so must be obtained. A subject of increasing concern is that of governments, notably that of the United States of America, entering the market for "0-day exploits" (Reuters, 2013) – newly discovered software vulnerabilities that enable remote compromise of machines. The existence and promotion of such a market is certainly a cause for concern (PDF) (Reporters Without Borders, 2013). The European Parliament, in late 2012, adopted a resolution against the export

trade in information technologies, particularly surveillance, that can be used to restrict or impinge on human rights (EU Parliament, 2012).

It becomes increasingly clear, therefore, that strict and independent audit, of the means of surveillance, surveillance requests, and data handling should be a key element of any proposed surveillance framework. Due to the inherent lack of transparency associated with these approaches this must be supported by stringent penalties for misuse of either powers or data, and supported by independent oversight.

5 SCOPE OF SURVEILLANCE

Focusing on internet-based surveillance of communications, many surveillance powers currently being debated or brought into force in Europe address, as has been noted above, two major approaches: increased powers to investigate the communications of targeted individuals, and increased blanket surveillance of large-scale population without suspicion. Both of these concepts, which are by no means mutually exclusive, bring their own risks. It is important to consider why these proposals are being made.

A key concept, certainly as expressed by the rhetoric that surrounds these measures, is the need to maintain capabilities to observe and analyse the communications of suspected criminals. The capability for wiretapping telephones already exists and, it has been repeatedly argued, unless this ability is extended to the internet then important communications will be beyond the reach of investigators. The recent NSA leaks have shown that this extension of scope to the internet is far more developed and widely deployed than had been thought.

As has been argued above, the telephone network cannot form a reasonable allegory for the internet in terms of surveillance policy. The number of daily activities and services, and the amount of data about individuals that is centralised in computers makes any access to this system inherently far more invasive than access simply to communications. Restricting the level of this access, so that reasonable levels of investigation into criminal activities can take place without violating the rights of internet users, is arguably the most important policy challenge in developing such schemes.

One important limitation of surveillance powers, that has rightly been the subject of much confusion, is the distinction between content data and communications data. This has often been expressed as the ability to see who is emailing whom, the communication, but without being able to read the content of those emails.

This intuitive explanation, which seeks to alleviate concerns of invasion of privacy, hides a great deal of complexity. While communications and content can easily be separated in a simple model such as email, the distinction is by no means so clear in other forms of internet traffic. For access to a website it is easy to see that the URL 'google.com' is distinct from the page of search results that you see when querying the site. In reality, however, a Google search for 'government surveillance' will result in a URL of the form 'http://www.google.com/search?q=government+surveillance'. With the search terms, and potentially other identifying information, so prominently highlighted, should this be considered communication or content data?

Within the UK, when these matters were first debated leading up to the Regulation of Investigatory Powers Act Great Britain, 2000, the UK's Foundation for Information Policy

Research successfully argued that only a URL up to its first forward slash, such as www.google.com/ should be considered communications data, with all other components of the URL being out of scope. With more and more complex services, the difficulty of separating communication from content data will continue to cause concern. Nor is the restriction to communications data a barrier against invasion of privacy. Whilst access to content is indisputably privacy-invasive, the patterns of communication can reveal a great deal about the nature of those communications. The frequency, duration, and length of messages can all reveal likely forms of conversation – it is precisely this type of information that analysts can exploit in the investigation of crime. If such information can be of use in investigating the patterns of communication between criminal suspects, it can equally be applied to communications between private individuals going about the course of their lives.

5.1 BUNDESTROJANER

In October 2011 the Chaos Computer Club, a German hacker group, published an analysis of software apparently installed on citizens' computers by various German police authorities for the purposes of spying on the activities of the operators (Chaos Computer Club, 2011). This software is now commonly known as the Bundestrojaner or Staatstrojaner ("Federal Trojan").

The German Federal Constitutional Court (*Bundesverfassungsgericht*) has ruled that the use of such software is permissible, and only permissible, for the purposes of wiretapping internet telephony, on the basis that wiretapping at the source of the transmission is the only way to gain access to the content of calls that are typically encrypted when transmitted across the internet. In order to restrict the capabilities of such approaches, any software to achieve these aims is required to be both technically and legally restricted to tapping of internet telephony. The Chaos Computer Club's analysis demonstrated that the software in question was capable of greatly expanded capabilities including the ability to log keystrokes, to enable a computer's microphone or internal camera. Most worryingly, the software allowed arbitrary control over the computer, including the ability to upgrade the trojan itself remotely, enabling arbitrary new functionality to be added once the software was initially installed.

5.2 LOPPSI2

While the German Staatstrojan received a significant level of attention in the German press when discovered, it is not an isolated example of such an approach in Europe. The passing, in 2011, of the French LOPPSI2 (*loi d'orientation et de programmation pour la performance de la sécurité intérieure*), a broad-based security law that includes a number of significant internet-based powers, enabled French law-enforcement and intelligence services to make use of similar approaches for the purposes of gaining access to internet-based communications.

The provisions in LOPPSI2 require authorisation of a magistrate before trojan software may be installed on a suspect's computer, through either local installation via physical access or remotely via the internet. Trojan access to a remote computer is permitted for four months in the first instance, with a single extension period. The use of such an approach is limited to 'the most severe cases'.

Despite the invasiveness of these approaches, they are still typically restricted to an individual's computer and not mandatory across an entire population for intelligence gathering purposes. As

has been discussed above, however, the potential abuses of such a system, providing almost total access and control over user devices, are extremely serious.

CONCLUSIONS

The surveillance measures discussed in this article seek to make society safer by providing a means to investigate criminal behaviour, either preventatively or otherwise. The wealth of data that the internet makes available provides a tempting set of inputs to the technical analyses that are increasingly available. At the same time, the rise of online services and encryption technologies at the heart of communications in modern Western societies begin to present challenges to the ability of law enforcement to investigate legitimate criminal suspects.

Despite this, it is crucial to consider not only the potential benefits of surveillance technologies, but to weigh them carefully against the serious systemic risks that such technologies enable. The powers of data mining techniques to invade privacy, when supplied with the detailed information about our lives captured by our online activities, are great.

While the ability to gain access to large volumes of information concerning targeted suspects is valuable, the ability of these techniques to identify suspicious behaviour in large populations is subject to inherent limitations. As such, the desire to gather, store, and analyse the data of entire populations requires significant justification that the risks of abuse, mission creep, security breaches and misapplication can be justified by their benefits.

The creation of the infrastructures required to gather such large amounts of data provide serious technical challenges, as well as opening the potential for abuse, either by third parties or, more likely, by misuse from within the bodies that have access to such data. With the inherent lack of transparency under which surveillance activities traditionally act, it is critical that any surveillance policy be subject to strict review and oversight. This approach and, more seriously, its limitations in terms of providing genuine accountability have been demonstrated in the FISA court system that oversees warrants for surveillance under the US Foreign Intelligence Surveillance Act. Whilst this body's approval is required, a Freedom of Information request by US Senator Harry Reid in April 2013 revealed that no warrants were denied by the court in the course of 1,789 requests in 2012 (U.S. Department of Justice, 2013).

Where surveillance seeks to operate at the level of service providers in a privatised landscape of web-based services, the cooperation of these service operators is necessary for access to data and, increasingly, encryption keys. The existence of an infrastructure required to enable such observation is already a great risk, and should both be avoided where possible or otherwise placed behind as many technical and institutional safeguards as possible. At the very least, no single authority should have control over such an infrastructure. The existence of PRISM and Tempora have suggested the existence of an institutionalised technical mechanism to surveil communications, functioning through streamlined surveillance requests made directly to service providers. A system of requests to third parties could allow for a level of oversight from companies themselves, allowing legal challenges against surveillance requests. In practice, it appears, this mechanism has not proven effective in the UK and the US.

Despite the recent NSA leaks, transparency, imposed both at a legal level and supported by the need to interact with private organisations that control infrastructure, remains one of the few potential mechanisms to mitigate the risks of abuse that inevitably accompany such approaches.

Where surveillance policy is based on targeted access to individuals' machines, as is seen in Germany and France, the inherent complications of gaining and maintaining access provide a useful limitation to the potential for wide-scale misapplication of the technologies. Despite this, such direct access to users' machines provides a truly frightening level of invasion into an individual's activities, and must be subject to severe limitations and audit.

To alleviate the concerns of broad-scale surveillance, any legal framework for enabling surveillance must, in the first instance, be based on the notion of targeted gathering of data on well-justified grounds. This precludes the *a priori* gathering and storage of data – such gathering should only occur in response to justified suspicion.

Where data is gathered and found not to be useful, particularly where it concerns third parties, it must be deleted as soon as it is determined not be of use for the narrow purpose for which it was initially gathered; ideally, this deletion would occur with some level of verifiability. There are also strong arguments to be made for the targets of such surveillance to be informed of the situation once any ongoing investigation is concluded.

The technological landscape in which we find ourselves is one in which the potential for surveillance is vast and growing. Institutions, understandably, wish to harness this potential but are predictably lax in considering its unavoidable limitations for preventing harm, and its serious risks of ongoing abuse and misapplication.

Surveillance law must therefore focus on restraining risks and abuses, without being carried away by false promises of effectiveness. Minimisation, decentralisation, accountability, targeted application, and limitation of access are all necessary steps to ensure that investigation of communications is a tool to protect and improve our society, rather than an ongoing source of harm and abuse.

FOOTNOTES

1. See R v Bowden [2000] 1 Cr App R 438, Archbold 31 - 108a.

2. In some states around the world, notably Iran, installation of a government-mandated digital certificate that allows a government to listen in on encrypted connections has, to some extent, been employed. Due to the extreme implications that this has for the general function of the internet, it is not considered here as a 'reasonable' option.

REFERENCES

Armitage, R. (2002). To CCTV or not to CCTV?. Retrieved from https://epic.org/privacy/surveillance/spotlight/0505/nacro02.pdf

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. Acta Psychologica 44: 211–233.

Base Rate Fallacy. (n.d.). In Wikipedia. Retrieved 26 May 2013 from https://en.wikipedia.org/wiki/Base_ rate_fallacy.

Chaos Computer Club. (2011). Chaos Computer Club analyzes government malware. Retrieved from: http://ccc.de/en/updates/2011/staatstrojaner

EU Parliament. (2012). A digital freedom strategy in EU foreign policy. Retrieved from

http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+TA+P7-TA-2012-0470+0+DOC+XML+V0//EN&language=EN

Reporters Without Borders. (2013). Enemies of the Internet - 2013 Report - Special Edition: Surveillance. Retrieved from http://www.reporter-ohne-grenzen.de/fileadmin/docs/enemies_of_the_internet_2013_01.pdf

Reuters. (2013). Special Report: U.S. cyberwar strategy stokes fear of blowback. Retrieved from http://www.reuters.com/article/2013/05/10/us-usa-cyberweapons-specialreport-idUSBRE9490EL20130510

United Nations Human Rights Council. (2013). Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression". Retrieved from http://www.ohchr.org/Documents/HRBodies/HRCouncil/RegularSession/Session23/A.HRC.23.40_EN.pdf

U.S. Department of Justice. Office of Legislative Affairs. Office of the Assistant Attorney General. (2013). FoIA Request Letter: US Department of Justice to Senator Harry Reid. Retrieved 31 July 2013 from http://www.wired.com/images_blogs/threatlevel/2013/05/fisacases.pdf

YouGov. (2012). Communications Data Bill. Retrieved from http://yougov.co.uk/news/2012/10/31/communications-data-bill/